Investment Memo

Eaton Corporation (NYSE: ETN)

Research period: August to October 2025. **Q3 2025 print:** November 4, 2025. **Position:** pass, no trade taken. **Prepared:** May 2026 as a thesis development memo.

1. The Framing

ETN was the first infrastructure-layer thesis I built. The research ran from August into October 2025, framing ETN as the cleanest grid-to-chip play in the AI infrastructure complex. The deeper output was conceptual. ETN was where the second-order bottleneck framework came together, and that framework later drove both the SMR pass and the CRWV entry.

The plan was to wait for the November 4 Q3 print before committing to any structure. On November 4, ETN reported record Q3 results alongside a $9.5B acquisition announcement that had landed the day before. The combination changed the company underneath the original thesis. I closed the file.

What this memo documents: the thesis as built, the framework evolution that produced it, the pass moment when the company shape changed, and how the tape played out over the following six months.

2. The Macro Thesis at the Time

The foundational insight was that AI is not just a software or semiconductor revolution. It is a physical infrastructure revolution. AI compute demand at the scale being committed by hyperscalers would massively increase electricity demand, strain existing grid infrastructure, and force utilities and data center operators into multi-year electrical capex cycles.

The logic chain was direct. AI growth requires more compute. More compute requires more electricity. More electricity requires more grid capacity, transmission, distribution, switchgear, transformers, and on-site power. The companies supplying those systems would see multi-year demand visibility that did not depend on which model architecture won or which hyperscaler captured the most workload.

The cleanest version of the thesis: regardless of which AI company wins, power infrastructure still needs to scale.

ETN sat at the center of that thesis. It was the picks-and-shovels exposure inside the AI buildout, with established cash flow and diversified industrial demand across utility, data center, aerospace, and commercial end markets. Not dependent on any single customer, GPU cycle, or AI architecture.

Three strategic moves had already played out by the time I was framing the thesis. The NVIDIA HVDC partnership announced in July 2025 positioned ETN as the certified electrical layer for next-generation AI data center designs. The Resilient Power Systems acquisition on August 6, 2025 added solid-state transformer technology to the portfolio, an architectural option that becomes load-bearing as power densities scale. PG&E's $73B utility capex plan through 2030, with ~10 GW of new demand visibility, provided multi-year pipeline cover for the broader electrical equipment cycle.

ETN's Q2 2025 print on August 5 delivered $2.95 in adjusted EPS, +8% YoY, with a record $7.0B in sales. But the company trimmed FY25 profit outlook on the same call. The stock sold off into mid-August. By late September, ETN was trading with a forward P/E near 28-30x, against diversified industrial peers at 20-26x.

3. Why ETN Specifically

The framework I was developing made ETN feel different from the speculative AI names. Lower upside, lower volatility, lower execution risk. Real revenue. Real customers. Real cash generation. Compared to CRWV, ETN had a fraction of the upside potential and a fraction of the binary outcome risk. Compared to SMR, ETN was a commercial business with a real customer base, not a pre-revenue commercialization story.

The deeper insight was about invisible winners. The biggest beneficiaries of a technological revolution are often the infrastructure providers, industrial suppliers, and bottleneck solvers, not the companies getting the most headlines. ETN was the test case for that framework. If the second-order bottleneck thesis was right, ETN would compound for years on a tailwind that did not depend on any single AI winner emerging. That was ecosystem exposure, not single-company exposure.

The framing was that the enabling layer might be safer and more durable than the application layer. Hype rotates. Infrastructure does not.

4. The Structural Concerns Identified Pre-Print

The thesis flagged five concerns going in. Each one became load-bearing by the time the Q3 print landed.

Valuation premium with sell-the-news sensitivity. ETN traded at the top of its peer range, forward P/E in the 26-31x band depending on source and day. At this multiple, earnings beats alone would not drive the stock. Guidance tone and revenue mix would.

AI capex slowdown risk. The thesis was indirectly dependent on sustained hyperscaler investment cycles. Any cooling in AI capex would compress infrastructure demand growth.

Vehicle and eMobility segment drag. The vehicle segment was a known weak spot. Soft North America truck and light vehicle markets were dragging on consolidated growth even as the data center business accelerated.

Industrial cyclicality. ETN was still an industrial company. Macro slowdowns, delayed utility spending, capex cuts, and slower construction cycles could all compress demand even if the secular AI infrastructure thesis was right.

Duration mismatch. The catalysts that would re-rate ETN materially sat on a multi-year timeline. The compounder thesis would take years to play out. Any short-dated option structure could not express the underlying conviction.

5. The Q3 2025 Catalyst That Crystallized the Pass

Eaton reported Q3 2025 on the morning of November 4, 2025. The headline numbers were strong:

- Adjusted EPS $3.07, up 8% YoY
- Record segment margins of 25.0%, above the high end of guidance
- Electrical Americas backlog up 20% YoY to $12B (all-time record)
- Data center orders accelerated 70% YoY; data center sales up 40%
- Book-to-bill 1.2 quarterly, 1.1 on a rolling twelve-month basis
- Full-year 2025 guidance reaffirmed at $11.97-$12.17 adjusted EPS

On paper, the print should have ratified the thesis. It did not. Three things in the disclosure changed the company underneath the original framing.

The Boyd Thermal acquisition, $9.5B, announced November 3. This was ETN's fourth acquisition of 2025. The deal added liquid cooling solutions to the data center portfolio at a 22.5x forward 2026 EBITDA multiple. Boyd was not expected to be EPS accretive until year two after closing, meaning 2027 at the earliest. The deal was strategically sound. It was also the largest acquisition in ETN's recent history, with integration complexity layered onto an already acquisition-heavy year (Fibrebond, Resilient Power Systems, Ultra PCS already closed).

The thesis I had built was "ETN is the pure-play grid-to-chip electrical layer for AI data centers." After November 3, the company was something different. It was now also a liquid cooling platform integrator with near-term EPS dilution, $9.5B in acquisition financing burden, and a multi-year integration overhang. The investment case had to be re-underwritten from scratch.

2026 capex guidance higher than 2025. Management flagged that capital expenditures would step up in 2026 to fund capacity expansion, then normalize in 2027. Combined with the Boyd financing, this meant free cash flow

pressure during the integration window. Premium-multiple stocks do not absorb FCF compression gracefully.

Stock reaction validated the valuation concern. ETN traded down ~5% on the Q3 print despite the record numbers. The pre-print analysis had specifically called this risk. At peer-high valuation, beats alone would not move the stock.

The framework I had built could not express the new thesis. The right move was to close the file, not force a structure to fit a company that had just changed shape.

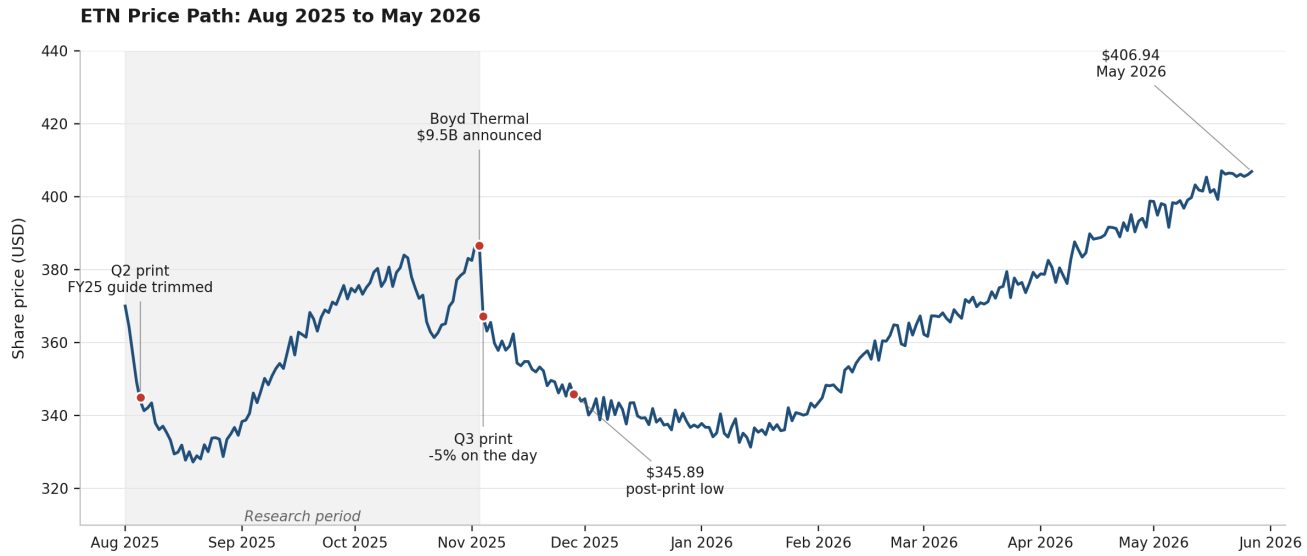

Figure 1. ETN closing prices with the key research-window events marked. Q2 guide trim, Boyd announcement, Q3 print reaction, and the six-month base before recovery into mid-2026.

6. How It Played Out

November 4 to November 28, 2025. Stock declined from the post-print level to ~$345.89 by November 28. Down approximately 10-11% over the month. The "sell the news on premium multiple" pattern the pre-print analysis had flagged was the dominant tape behavior.

Late November 2025. Olivier Leonetti, EVP and CFO, announced his departure effective April 1, 2026. CFO transitions during major M&A integration windows are governance concerns by definition.

December 2025 through Q1 2026. Analyst price targets diverged. RBC Capital raised to $432 on the data center order strength. Other coverage flagged integration risk and premium-multiple sensitivity. The stock traded in a wide range as the market argued with itself about whether ETN was a strategic-scarce-asset buyer or a late-cycle over-payer.

Current state (May 2026). ETN trading around $407, with a $450.60 consensus 12-month price target across 30 analysts. Morgan Stanley raised to $500 (Overweight). Evercore raised to $453 (In Line). Full-year 2025 results delivered $27.45B in revenue (+10.33%) and $4.09B in earnings (+7.72%). The business is performing. The stock has spent six months working off the post-print compression.

The pass did not generate a P/L event. It generated capital that did not get tied up in a position whose original thesis no longer described the company.

7. What This Research Produced

The ETN work generated more lasting output as a framework than as a trade. Four insights came out of it that I have used in every subsequent memo.

Second-order bottleneck investing. The core question was: what physically breaks if AI actually scales? The answer in 2025 was electrical infrastructure, grid capacity, power management. That framework directly drove the

SMR research, where the same logic extended to nuclear baseload as a deeper second-order bottleneck. It also informed the CRWV entry, where compute capacity itself was the bottleneck being priced. The dependency chain of AI growth into compute demand into electricity demand into grid capacity into infrastructure spending became one of the most useful analytical structures I built across the research practice.

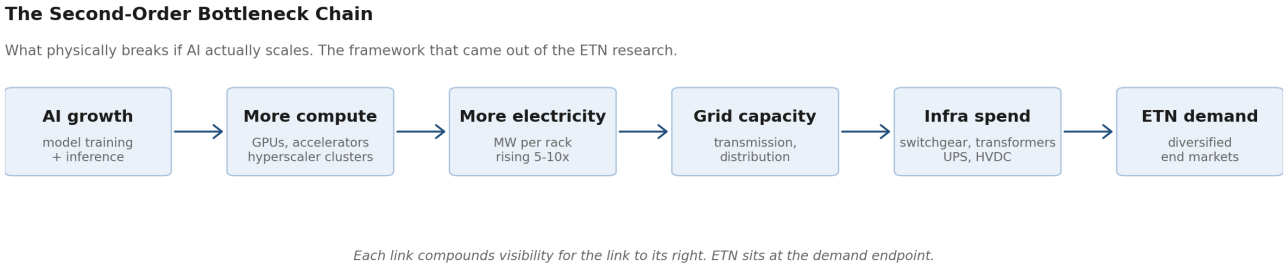

Figure 2. The second-order bottleneck chain. Each link compounds visibility for the link to its right.

Invisible winners. The realization that infrastructure providers and bottleneck solvers often capture more of a technological revolution's value than the visible end-product names. This reframed how I evaluated AI-related ideas going forward. The question shifted from "which AI company wins" to "what gets bought regardless of which AI company wins." That reframe was what made the CRWV entry possible weeks later.

Ecosystem exposure over single-company exposure. ETN's diversified end markets and the picks-and-shovels framing produced a different way of thinking about AI exposure. Single-name AI bets carry platform risk, model risk, and customer-concentration risk. Infrastructure-layer exposure spreads those risks across the entire buildout. The enabling layer is often more durable than the application layer because hype rotates and infrastructure does not.

Acquisition-as-thesis-change risk. A new framework dimension that came out of the Q3 print. The Boyd deal demonstrated that management at high-quality industrials can make company-redefining capital allocation decisions inside a planned option duration. For premium-multiple names with active M&A programs, "the company you researched might not be the company that reports earnings" is a real risk vector. The framework now has a slot for tracking M&A propensity in the months leading up to a planned trade entry, and a slot for cumulative integration burden across acquisition-heavy operators.

The pass was the right call. The framework that produced it became foundational to everything that followed.

Personal investment research. Not investment advice. Not affiliated with any registered investment firm. Sources: Eaton Corporation investor relations and SEC filings (Q2 2025, Q3 2025), earnings call transcripts, NVIDIA partnership disclosures, Boyd Thermal acquisition announcement (November 3, 2025), public market data, and analyst coverage from RBC Capital, Morgan Stanley, and Evercore ISI.